

Adaptive sampling for rover x-ray lithochemistry

David R. Thompson, Abigail Allwood, Christopher Assad, David Flannery, Robert Hodyss, Emily Knowles, Lawrence Wade

Jet Propulsion Laboratory, California Institute of Technology
e-mail: {firstname.lastname}@jpl.nasa.gov

Abstract

Planetary rovers arriving at a new site may only have time to deploy contact instruments at a few selected locations before the mission must progress to the next locale. Since the number of command cycles is limited, operators typically script entire measurement sequences in advance. Onboard data analysis can improve the quality of these measurements by reacting immediately to collected data and optimizing the measurement sequence on the fly. Here we present a method by which a smart arm-mounted contact instrument can improve the efficiency of transect and raster patterns, selectively sampling from distinctive materials instead of distributing its time budget evenly across repetitive points. We focus on the specific case of contact measurements for x-ray fluorescence spectroscopy. Simulated and laboratory experiments demonstrate significant improvements in science data yield.

1 Introduction

Planetary surface survey missions generally face challenging time constraints. Rovers with contact instruments may only be able to measure a few targets at each site before the mission must progress to its next locale. On Earth, field geologists faced with time pressure would use fast exploratory measurements to find the important targets and focus their study. However, this is not always possible for planetary rovers; the number of command cycles at each site is limited, and entire measurement sequences are often scripted in advance. This can lead to redundancies or gaps in data collection, particularly for instruments with long integration times such as the Alpha Particle X-Ray Spectrometer (APXS) [Rieder et al., 2003] or Dynamic Albedo of Neutrons (DAN) [Litvak et al. 2008]. Without knowing which features will be most interesting, missions must compromise and use shorter integration times over more sample sites, reducing sensitivity. If deployed in transect or raster patterns, the sampled area must provide a margin around the targets of interest to accommodate uncertain instrument pointing. Moreover, there is a trend of higher spatial resolution

measurements that co-locate compositional data with small-scale textures, fabrics, and structures in the rock [Wade et al. 2012]. High spatial resolution provides valuable contextual information for interpreting potential biosignatures, but makes efficient sample allocation even more critical.

In these scenarios, instrument autonomy can improve data quality by reacting immediately to data and adapting the measurement sequence [Smith et al. 2007, Thompson et al. 2013]. This paper describes a method for onboard data analysis that can immediately recognize distinctive units from instrument data. It allocates integration time to favor a representative sampling rather than evenly distributing observation time across many repetitive points. This can achieve a significant improvement in time efficiency while approaching the fidelity and science value of exhaustive measurements.

We focus our study on contact instruments, with a case example of x-ray fluorescence spectroscopy [Hodyss et al. 2012, Wade et al. 2012]. We consider the Planetary Instrument for X-Ray Lithochemistry (PIXL) that demonstrates a well-defined tradeoff between sampling time and measurement accuracy. PIXL is a rover arm-mounted spectrometer proposed for future rover missions such as a Mars 2020 rover. Placed next to the target, it excites the surface with x-rays producing a fluorescence spectrum that indicates elemental composition. Uniquely among flight instruments, PIXL incorporates a source-ray focusing optic with very high spatial resolution. Rovers can deploy it with multiple acquisitions in transects or grid patterns to generate fine scale elemental maps. Integration times can vary from a few seconds to a minute or longer. One second is sufficient for a quick scan of major elements [Hodyss et al. 2012], while longer integrations provide accurate compositional information for trace element analysis.

Our adaptive data collection approach is based on unsupervised machine learning principles that recognize distinctive materials and modify the integration time. This provides a survey map indicating major compositional units, with at least one authoritative long integration for each type. We have implemented the algorithm in the instrument control software for the PIXL breadboard instrument, where it has proven effective in laboratory demonstrations.

We begin with a discussion of the instrument and an adaptive sampling methodology based on principles of vector quantization. We then describe the experimental approach and results. Tests on heterogeneous “worst case” samples have achieved 60% reductions in the number of long integrations. We find an even larger savings for homogeneous samples. This simple approach is easy to implement in embedded processors and applies to a wide range of planetary exploration instruments. We will continue to test this method as a component of the control strategy. We conclude with a discussion of other instrument platforms that could benefit from this approach, and of future development.

2 Micro-XRF instrument and datasets

PIXL is a high spatial resolution X-ray fluorescence (Micro-XRF) instrument that measures a rock’s elemental composition by exciting the target with a known source and analyzing the resulting fluorescence. Its spatial resolution of 100 microns can analyze features as small as individual laminae or sand grains [Wade et al. 2012]. The physical instrument consists of an emitter/detector pair placed within a few millimeters of the target. It is actuated to translate the field of view horizontally. The entire instrument can mount on a motorized end-effector providing a second degree of freedom, enabling full 2D mapping.

Figure 1 shows two typical spectra from the instrument. Peaks at different energy levels indicate the relative abundance of different bulk and trace elements. The precise relationship between peak height and elemental composition is subtle, but broadly speaking a change in peak height indicates a measurable difference in the material being measured. As the integration time increases, the number of counts in each energy level grows causing the peaks to become more distinct and reducing noise. The peaks can be modeled with high precision, enabling very accurate determination of trace elemental compositions.

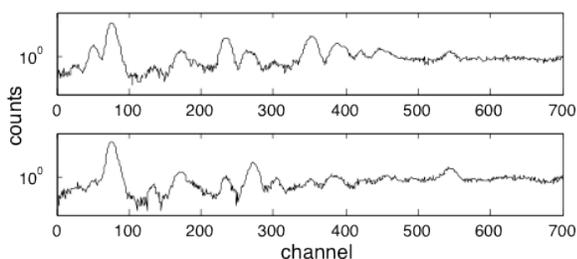


Figure 1: Example spectra acquired from two different materials in the same sample (log counts are shown).

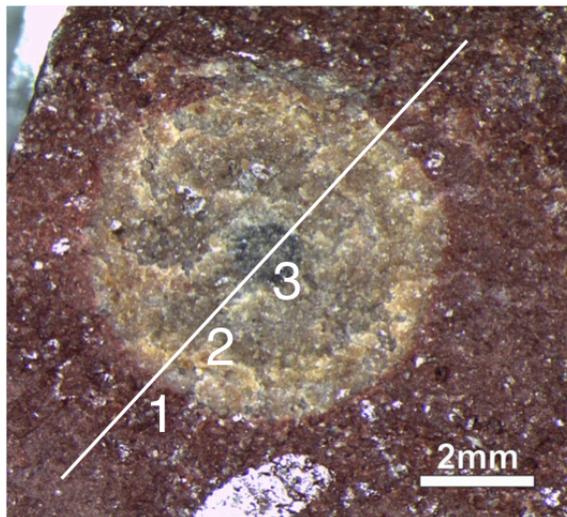


Figure 2: Reduction spot image and transect.

We simulated adaptive sampling performance with two samples of different complexity. Figure 2 shows the first dataset that was collected from an altered Neoproterozoic sandstone [Parnell 2014]. The concentric rings of alteration are visible in the center of the image as a “reduction spot”, a potential microbial biomarker [Spinks et al. 2010]. We acquired a sequence of 130 x-ray spectra indicated in the image by the white line. It transected at least three zones: (1) the unaltered red (ferric) sandstone, (2) the light-colored outer zone of alteration, and (3) the dark-colored, Vanadium- and Copper-rich interior of the reduction spot.

Figure 3 shows a more complex sample: a polymictic conglomerate consisting of many different materials in close spatial proximity. We analyzed a vertical transect of this “mud conglomerate”, indicated here by the white vertical line. This transect has 180 spectra. In the image, four numbers indicate distinctive surface features: (1) a silicified sandstone conglomerate with diverse clasts; (2) a silicified tuffaceous mudstone containing fuchsite, (3) a heterogeneous vein structure, enriched in Titanium and other elements; and (4) a Fe/Mn carbonate alteration. Figure 4 shows two elemental maps resulting from a complete raster of the sample at full spatial resolution. The sample surface is heterogeneous, with isolated reduction spots having higher iron and titanium concentrations. The maps correlate these compositional variations with meso-scale textural differences visible in sample. While such products are powerful tools for planetary exploration, they require thousands of integrations at up to a minute per location; this limits their use to selected sites of greatest scientific interest.

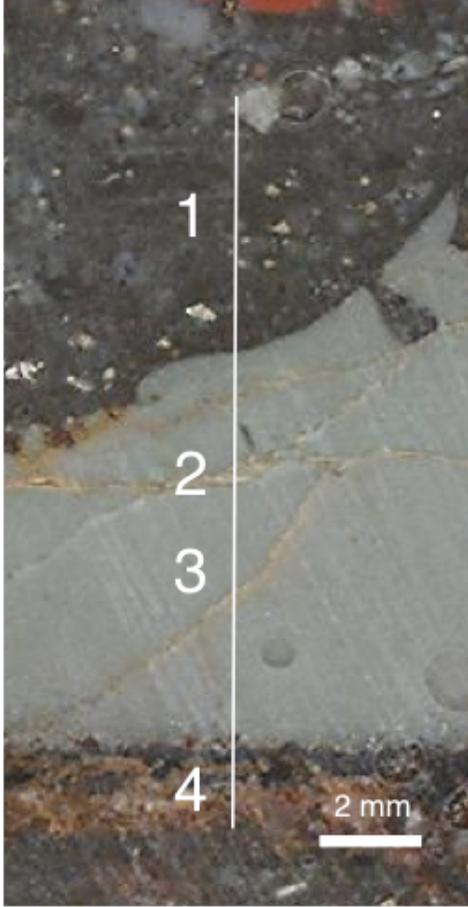


Figure 3: “Mud conglomerate” image and transect.

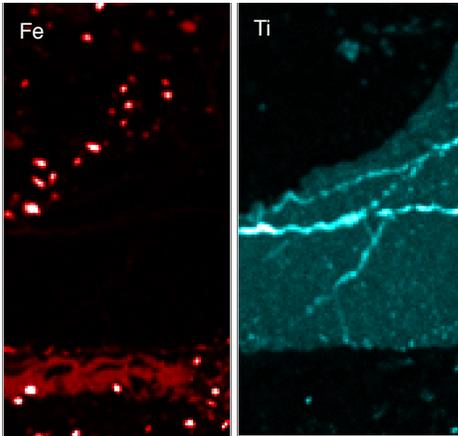


Figure 4: Elemental maps of Fe and Ti from the mud conglomerate sample

3 Adaptive sampling method

This section describes an adaptive sampling approach that can provide high spectral fidelity at a few selected locations without the time cost of an exhaustive map. Our method rests on a simple but intuitive distinction between *rough classification* (for which a short integration is sufficient) and *accurate compositional analysis* (which requires a long integration for optimal SNR). We assume spectral datapoints are presented one at a time to the instrument. With each new datapoint, an irrevocable decision must be made about whether to linger for a long integration or to proceed to the next location. The objective is that all survey points be sufficiently similar to at least one representative long dwell spectrum, with similarity defined according to a meaningful spectral distance measure. This ensures sufficient data to perform a full compositional analysis for each distinct material.

We begin the analysis by transforming the spectrum into a vector of elemental peak heights, $a \in \mathbb{R}^p$. Each entry represents the sum of detector counts over a relevant interval, normalized to the silicon band. The normalization scales values to manageable ranges, controls for variable integration times, and also accounts for minor changes in sensitivity across acquisitions. Silicon is typically the dominant element, so this normalization also brings the entire vector closer to a measurement of fractional composition.

Each transect consists mainly of short “survey” points $A = \{a_1 \dots a_n\}$, but some subset of these will also be “long dwell” points $B = \{b_1 \dots b_m\}$ $B \subseteq A$. Selection of long dwell points is tantamount to a vector quantization problem from signal compression, in which a small number of complete template vectors b stand in for the complete signal elsewhere. We define a distance measure as a simple Euclidean metric with an isotropic rescaling of dimensions by a vector w of nonnegative length scales. This yields the following global cost function:

$$L(B) = \sup_a (\min_{x \in B} |w^T a - w^T x|_2) + C(B) \quad (1)$$

This expression has two terms. The first represents the maximum distortion: the maximum spectral distance between any collected datapoint and its most similar long dwell spectrum. The term on the right is a cost function that increases monotonically with the number of long dwell datapoints B . We use a simple iterative algorithm that dates at least to the early vector quantization studies of [Paul 1983]. The instrument maintains a library of all long dwell spectra collected during the transect. At each new survey point, it measures the weighted distances between the new spectrum a and each library

spectrum b . Distances exceeding a predefined threshold T trigger a new long integration. In this way, the library grows to ensure that the left term of Equation 1 remains bounded. If collecting a new long integration would cause the instrument to exceed its total time allocation for the transect, it forgoes all further long integrations and finishes the remainder of its sequence. Algorithm 1 below shows the entire procedure. It requires only a handful of floating point operations and is suitable for an embedded instrument processor.

Input: peak weight coefficients $w \in \mathbb{R}^n$,
total budget M for long integrations,
distance threshold T

Output: survey spectra A ,
long dwell spectra B

$A, B \leftarrow \{\}$

While the transect is not finished

acquire a short survey spectrum a

$A \leftarrow A \cup a$

If $M > |B|$ and $\min_{x \in B} |w^T a - w^T x|_2 > T$

acquire a long integration spectrum b

$B \leftarrow B \cup b$

Algorithm 1: Online adaptive sampling

Operators can alter the system behavior in several ways. First, they can change T to adjust the rate of data collection. This value determines the tolerance for triggering long integrations. Second, they can impose a time budget M as a final check to ensure that the system does not exceed its total resource allotment under any circumstances. Finally, operators can alter the weight vector w in order to emphasize or de-emphasize specific elements. In practice it might be simpler to set channel weighting factors using a representative training set of similar library samples or prior measurements from a nearby locale. For this work we compile at least 50 spectra from a prior dataset, and use the inverse standard deviations of the vector elements in that population as weighting coefficients. This ensures that all the values are scaled to an appropriate order of magnitude.

Note that the adaptive algorithm triggers based on a distance threshold; this has the effect of bounding the maximum distortion. There are alternative objective functions from the vector quantization literature such as the mean square distortion, which might favor different trigger criteria. Additionally, other distance measures are possible such as a full covariance matrix in lieu of isotropic weights. That would entail a Mahalanobis distance formulation. We refer the reader to more canonical texts on vector quantization [Makhoul et al.

1985] for a comprehensive discussion of these alternatives.

Figure 5 illustrates this process visually for the mud conglomerate sample, using just two channels: an iron channel and a titanium channel. Most datapoints crowd near the origin, but there are also sparse clouds of outliers having elevated levels of each element. The cloud at the top of the plot corresponds to region (4) of Figure 4. The two points at right lie on the titanium enriched veins of region (2). Here for illustrative purposes we use a strict threshold causing the adaptive sampling to select just four points. It selects follow up points that are evenly spaced within the 2D space defined by these axes. In this way both veins and iron alterations are sampled preferentially.

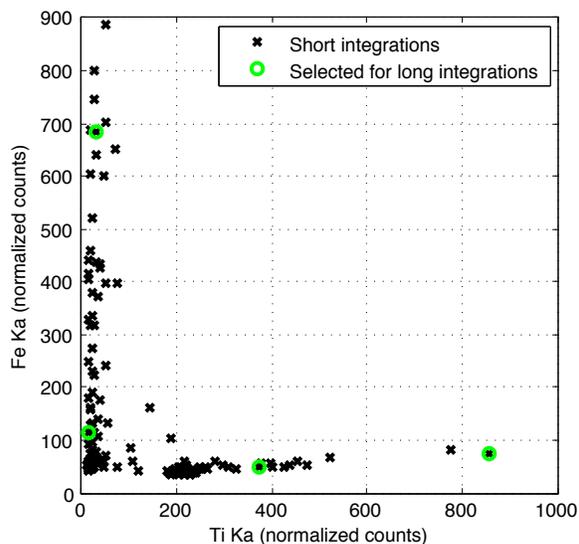


Figure 5: Adaptive sampling of the mud conglomerate sample, illustrated for just two elements.

4 Experimental results

We applied adaptive and evenly-spaced sampling methods to the reduction spot and mud conglomerate datasets. We simulated data collection by presenting each datum sequentially to a virtual instrument, and recording those selected for followup. Figure 6 shows simulated performance for the reduction spot sample, in terms of the maximum distortion (the largest weighted Euclidean distance from a short spectrum to its nearest long integration). We calculated distortion scores with the same weighting coefficients used for sampling. The image shows simulation trials as red and blue circles, with trend lines given by locally linear smoothing. For

very dense transects with many samples, an evenly spaced transect approaches the performance of adaptive sampling. However, for a fixed number of spectra, distortions using adaptive sampling are typically half those of evenly spaced spacing. It takes less than 20 spectra for adaptive sampling to achieve the fidelity of 100 evenly spaced samples.

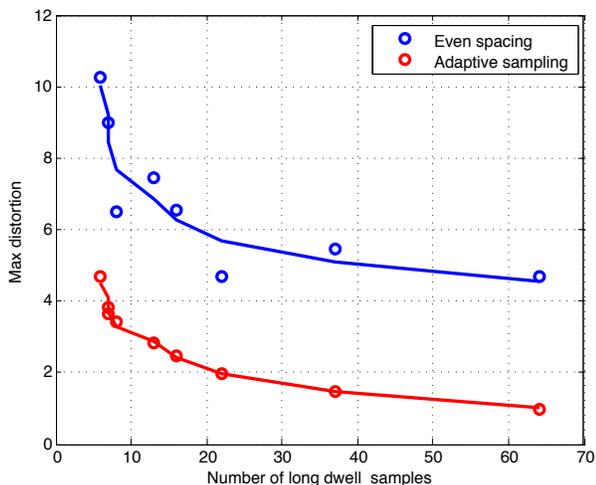


Figure 6 Simulated performance on the reduction spot dataset.

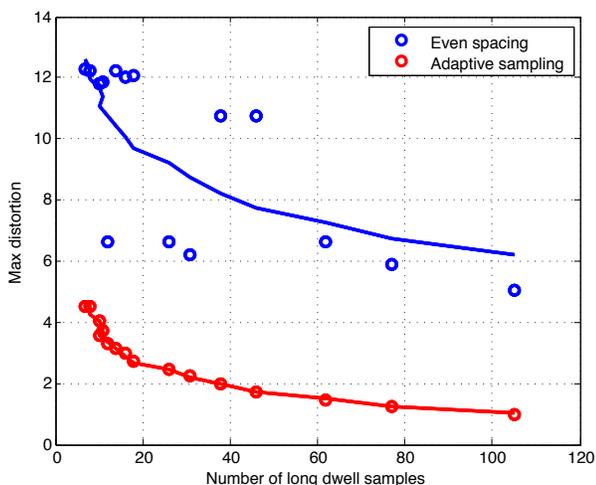


Figure 7: Simulated performance on mud conglomerate dataset.

Figure 7 shows the performance for the mud conglomerate sample. Here adaptive sampling yields more significant benefits; just 6 long integrations provide better fidelity than the even sampling scheme achieves with over 100. The adaptive sampling is also

considerably more consistent. This is due to the presence of isolated structures - particularly veins, but also the iron alteration - that appear in just one or two locations and are difficult to hit by chance. Any such unit that is ignored dominates the maximum distortion score, but the adaptive sampling typically finds them all.

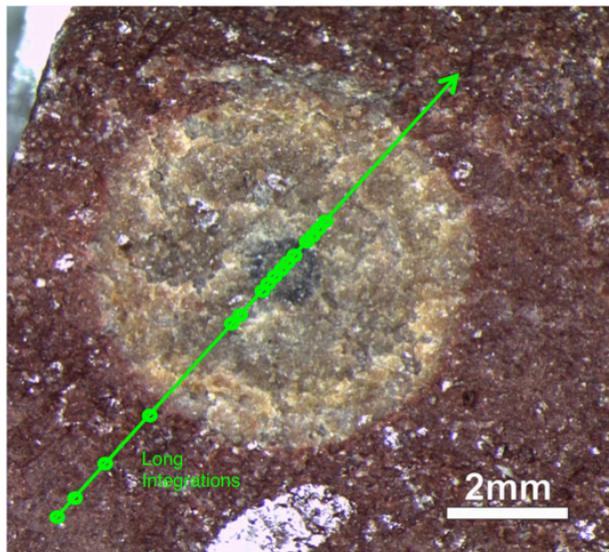


Figure 8: Real time adaptive sampling result from the laboratory breadboard instrument.

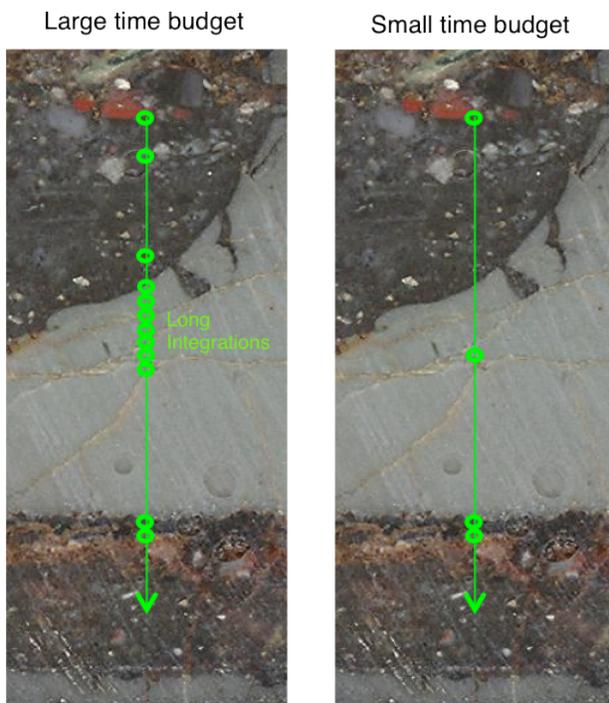


Figure 9: Real time adaptive sampling result from the laboratory breadboard.

We also tested the adaptive control system in real time demonstrations using the PIXL breadboard instrument. We programmed Algorithm 1 into the control system and performed an adaptive transect of the reduction spot sample using 130 short survey spectra. We set weighting coefficients using standard deviations of a prior dataset from the same rock. Figure 8 shows the result: the adaptive system triggers a long integration 17 times out of 130 opportunities. There are a few initial triggers at the start of the transect, and some more the border of the light-colored alteration zone. There is also a series of many consecutive triggers as the instrument crosses the inner alteration zone and a heterogeneous region on the opposite side. After 75 spectra, the sample is adequately characterized and the system never performs another long integration. This threshold achieves an 87% savings in long integrations vis a vis exhaustive sampling.

We also evaluated performance for the more challenging mud conglomerate sample. Here we tried two different thresholds to represent lenient and strict time constraints. Each transect contained 40 short survey spectra. Figure 9 shows the result. The lenient threshold triggered a long integration 12 times, for a 70% reduction in long integrations, while still acquiring samples from all four units. When using the strict threshold it triggered just 4 times for a dramatic 90% reduction, but it still acquired a long integration spectrum from all units except the silicified mudstone. It is possible that revising to the weighting vectors would recover samples from the fourth unit. It is likely that more redundant sample patterns (such as denser transects or 2D maps) would glean additional time savings.

5 Conclusions

This work describes a method for adaptive sampling by an x-ray fluorescence instrument and reports on successful real time tests in a laboratory setting. There are several obvious paths to refine these initial proofs of concept. Potential enhancements include the use of learned Mahalanobis distance metrics to account for correlations in elemental compositions [Francis et al., 2014]. Moving beyond spectral analysis, an even more powerful approach might be to target the instrument based on onboard analysis of the context image. The visual analysis could involve mapping specific visual surfaces as in [Bekker et al., 2014] or [Wagstaff et al., 2013], or segmenting the image and identifying local anomalies as in [Thompson et al., 2013, Francis et al., 2014]. While there is scope for further development, we believe the existing algorithm strikes a good balance

of simplicity and flexibility to accommodate different measurement objectives. It provides a palette of behaviors that operators can specify at command time to improve the efficiency and science yield of X-ray fluorescence spectrometers. More generally, the method applies equally well to other operations scenarios, such as a multi-instrument system that uses a fast survey sensor to selectively deploy a second more costly measurement [Chien et al., 2014]. Adaptive instrumentation will be increasingly important as future missions continue to improve collected data volumes despite persistent limits on bandwidth and latency.

6 Acknowledgements

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration. We thank John Parnell of the University of Aberdeen for the “reduction spot” sample. We acknowledge the support of the NASA Astrobiology Science and Technology Instrument Development program (NNH10ZDA001N-ASTID) and NASA AMMOS Technology Development. Copyright 2014 California Institute of Technology. All Rights Reserved.

References

- [Bekker et al. 2014] D. Bekker, D. R. Thompson, W. Abbey, A. Allwood, N. A. Cabrol, T. Fuchs, K. Wagstaff. A field demonstration of a smart instrument for automatic classification of geologic surfaces. *Astrobiology* (in press).
- [Chien et al. 2014] S. Chien, B. Bue, J. Castillo-Rogez, D. Gharibian, R. Knight, S. Schaffer, D. R. Thompson, K. L. Wagstaff (2014). Agile Science: Using Onboard Autonomy for Primitive Bodies and Deep Space Exploration. SpaceOps 2014.
- [Francis et al. 2014] R. Francis, K. McIsaac, D. R. Thompson, G. R. Osinski, Autonomous mapping of outcrops using Multiclass Linear Discriminant Analysis. *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Montreal, Canada, 2014.
- [Hodyss et al. 2012] R. Hodyss, L. Wade. Micro-XRF: Elemental Analysis for in Situ Geology and Astrobiology Exploration (2012). JPL Technical Report. <http://hdl.handle.net/2014/43342>
- [Litvak et al. 2008] Litvak, M.L.; Mitrofanov, I.G.;

- Barmakov, Yu.N.; Behar, A.; Bitulev, A.; Bobrovniksky, Yu.; Bogolubov, E.P.; Boynton, W.V. et al. (2008). "The Dynamic Albedo of Neutrons (DAN) Experiment for NASA's 2009 Mars Science Laboratory". *Astrobiology* 8 (3): 605–12.
- [Makhoul et al. 1985] Makhoul, J., Roucos, S., & Gish, H. (1985). Vector quantization in speech coding. *Proceedings of the IEEE*, 73(11), 1551-1588.
- [Paul 1983] Paul, D. B. (1983). An 800 bps adaptive vector quantization vocoder using a perceptual distance measure. In *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP'83*. (Vol. 8, pp. 73-76). IEEE.
- [Parnell 2014] Parnell, John (2014). University of Aberdeen Dept. of Geology and Petroleum. Personal exchange.
- [Rieder et al. 2003] R. Rieder, R. Gellert, J. Brückner, G. Klingelhöfer, G. Dreibus, A. Yen, S. W. Squyres (2003). "The new Athena alpha particle X-ray spectrometer for the Mars Exploration Rovers". *J. Geophysical Research* 108 (E12): 8066.
- [Smith et al. 2007] Smith, T., Thompson, D. R., Wettergreen, D. S., Cabrol, N. A., Warren-Rhodes, K. A., & Weinstein, S. J. (2007). Life in the Atacama: Science autonomy for improving data quality. *Journal of Geophysical Research: Biogeosciences (2005–2012)*, 112(G4).
- [Spinks et al. 2010] Spinks, Samuel C., John Parnell, and Stephen A. Bowden (2010). "Reduction spots in the Mesoproterozoic age: implications for life in the early terrestrial record." *International Journal of Astrobiology* 9 (4): 209.
- [Thompson et al. 2013] Thompson, D. R., Cabrol, N. A., Furlong, M., Hardgrove, C., Low, B. K. H., Moersch, J., & Wettergreen, D. (2013, May). Adaptive sensing of time series with application to remote exploration. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on* (pp. 3463-3468). IEEE.
- [Wade et al. 2012] Lawrence A. Wade, Robert Hodyss, Abigail Allwood, Ning Gao, and Kris Kozaczek (2012). A novel Micro-XRF for in-situ geological exploration of other planets. Denver X-Ray Conference, 2012.
- [Wagstaff et al. 2013]. K. L. Wagstaff, D. R. Thompson, W. Abbey, A. Allwood, D. L. Bekker, N. A. Cabrol, and T. Fuchs. (2013). Smart instrument classification for in situ rock and layer analysis with TextureCam. *Geophysical Research Letters* Volume 40, Issue 16, pages 4188–4193, 28 August 2013.